


REPORT DOCUMENTATION PAGE

Public reporting burden for this collection of information is estimated to average 1 hour per response, including gathering and maintaining the data needed, and completing and reviewing the collection of information, collection of information, including suggestions for reducing this burden, to Washington Headquarters Service, Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paper

0563

1. Sources
of this
Jefferson

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE		3. REPORT TYPE AND DATES COVERED FINAL REPORT 01 Sep 94 - 30 Apr 95	
4. TITLE AND SUBTITLE High-Performance and Low-Cost Optical Interconnects				5. FUNDING NUMBERS 63218C 1601/10	
6. AUTHOR(S) Dr John A. Neff					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Colorado Campus Box B-19 Boulder, CO 80309-0019				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR/NE 110 Duncan Avenue Suite B115 Bolling AFB DC 20332-0001				10. SPONSORING/MONITORING AGENCY REPORT NUMBER F49620-94-1-0443	
11. SUPPLEMENTARY NOTES					
12a. DISTRIBUTION/AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE: DISTRIBUTION UNLIMITED				12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) <i>This grant developed models of optoelectronic technology to predict manufacturing cost for monolithic and hybrid integration. For the hybrid CMOS-SEED technology, the models predict that modulator yield limits the overall system yield. In addition, the models show that CMOS-SEED is lower cost than the monolithic FET-SEED, though much more expensive than conventional silicon due to the high GaAs epitaxial wafer cost. Hence, CMOS-SEED integration will be limited to small chips on MCMs. The yield models also predict the ratio of optoelectronic interconnects to transistors in a balanced system. In addition, for the same cost systems, we showed the performance, reliability or architectural advantage necessary to make optoelectronic interconnects competitive with electronics.</i>					
14. SUBJECT TERMS				15. NUMBER OF PAGES	
16. PRICE CODE				17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	
18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED		19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED		20. LIMITATION OF ABSTRACT UNCLASSIFIED	



DTIC QUALITY INSPECTED 8

19951002 017

High-Performance and Low-Cost Optical Interconnects:

AFOSR Final Report

Charles Stirk and John Neff
University of Colorado
Boulder, CO 80309-0525
charlie@boulder.colorado.edu
303-492-2896 fax: 3674

CU Account No: 1536846
DoD Contract No: DOD F49620-94-1-0443

Abstract

This grant developed models of optoelectronic technology to predict manufacturing cost for monolithic and hybrid integration. For the hybrid CMOS-SEED technology, the models predict that modulator yield limits the overall system yield. In addition, the models show that CMOS-SEED is lower cost than the monolithic FET-SEED, though much more expensive than conventional silicon due to the high GaAs epitaxial wafer cost. Hence, CMOS-SEED integration will be limited to small chips on MCMs. The yield models also predict the ratio of optoelectronic interconnects to transistors in a balanced system. In addition, for the same cost systems, we showed the performance, reliability or architectural advantage necessary to make optoelectronic interconnects competitive with electronics.

Optoelectronic Manufacturing Cost

Under our AFOSR grant, we constructed manufacturing cost models for leading optoelectronic integration technologies. This grant is a follow-on to our earlier AFOSR sponsored research where we analyzed the cost of monolithic FET-SEED and silicon processes. In the present grant, we extended the analysis to CMOS-SEED based systems.

The AT&T CMOS-SEED technology uses commercial silicon CMOS with optical modulators/detectors solder bumped to pads on the silicon. The modulators are 20 by 50 microns on a side. A typical CMOS-SEED system is a 2D smart pixel array with individual pixels consisting of 2x1 switches on a 250 micron pitch. Each pixel uses four modulators and eight solder bumps.

The yield of a process step is the probability that a component has no fatal defects. The modulator yield is limited by MBE spitting defects at 100/cm² density. The solder bump yield is one failure every 10⁵ bumps. The silicon yield is from the Semiconductor Industry Association Roadmap for 1995 with a density of 0.43 defects/cm². Figure 1 shows that in our typical system, the yield is limited by the multiple quantum well (MQW) modulators.




Figure 1. Yield as a function of chip area for the individual steps of silicon fabrication, solder bump attachment, and MQW modulator growth.

Besides the yield, the other important parts of the cost model are the material and process costs. The MQW wafers cost \$ 5,000 for a two inch wafer. On the other hand, a commercial silicon wafer costs \$ 1,250 fully processed. We combined the yields and the costs to create a model of the CMOS-SEED manufacturing sequence.

As shown in Fig. 2, the high MQW wafer cost makes the resulting smart pixel chips very expensive. In contrast, processed CMOS costs around \$ 4/cm². Figure 2 also shows the exponentially growing cost as a function of chip area which comes from yield loss. The exponential rises very quickly due to the high MQW defect density.

Solder Modulator Cost




Figure 2. Cost as a function of chip area for the CMOS-SEED smart pixel array.

Under the previous grant from AFOSR, we considered the monolithic FET-SEED process which can make systems functionally equivalent to those with the hybrid CMOS-SEED process. The cost advantage of the FET-SEED process is that there is no extra yield loss step due to hybrid solder bumps as in the CMOS-SEED. However, the FET-SEED transistors are susceptible to the MQW wafer defects, which have a density several orders of magnitude higher than those in silicon processes. Figure 3 shows that the monolithic transistor yield penalty is dramatic, and thus, the hybrid CMOS-SEED is much lower cost. Whenever practical, CMOS-SEED should be used for low-cost systems instead of the monolithic FET-SEED.

For	
I	<input checked="" type="checkbox"/>
d	<input type="checkbox"/>
ion	<input type="checkbox"/>
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	




Figure 3. Cost as a function of chip area for identical architecture CMOS-SEED and FET-SEED smart pixel arrays.

In the previous report, we showed that the FET-SEED smart pixel arrays with small numbers of pixels would be lower cost if they were composed of several chips that were solder bumped to a common substrate like an MCM. Figure 4 shows that for larger arrays, this is also true for the CMOS-SEED process. In particular, for a 256 switch array the lowest cost implementation is partitioned into two to four chips.

256 Node Module Cost




Figure 4. Cost as a function of chip area for the CMOS-SEED smart pixel array.

In summary, the hybrid CMOS-SEED manufacturing process is lower cost than the monolithic FET-SEED. However, CMOS-SEED is still much more expensive than silicon due to the high MQW wafer costs. In addition, CMOS-SEED will be limited to small chip sizes in an MCM because of the high MQW defect density.

Cost and Performance Tradeoffs in Optical Interconnects

In the last report, to integrate cost and performance we compared the implementation of optoelectronic multiplexers and demultiplexers in terms of latency, area (cost), simultaneous switching noise and power dissipation. In the follow-on grant, we investigated cost-performance tradeoffs based on yield.

The many performance comparisons between electrical and optical interconnects in terms of power dissipation, skew, and density largely neglect cost. The resulting system

demonstrations perform well, but are too costly to become products. This study takes a different approach that compares the performance of systems that cost the same. To force the cost to be the same, the relative number of each type of device in the system depends on the device manufacturing defect rates. The device types included in the comparison are CMOS transistors, solder bumps and wire bonds for electronics, and monolithic GaAs or hybrid on silicon for optoelectronics.

The following table illustrates the present defect rates in interconnect and logic technology. The defect density is the raw density of fatal defects in a given technology. Size accounts for the fact that devices have different sizes that are susceptible to point defects. Normalizing the raw density by the device size produces a defect rate that is a function only of the device type.

	Defect Density Size [μm^2]		Device Rate
CMOS Transistor	0.3/cm ²	10	10 ⁻⁸
Solder Bump	10 ⁻⁵ /bump	-	10 ⁻⁵
Wire Bond	10 ⁻³ /bond	-	10 ⁻³
Epitaxial GaAs	100/cm ²	100	10 ⁻⁴
		10	10 ⁻⁵

Balanced system design devotes limited resources to parts of a system in an attempt to optimize some system metric. Balancing denotes the change in subsystem contributions to the metric as the resources are shifted.

Consider balancing the manufacturing costs of the present microprocessor architectures. To equalize the yield per step, the device defect rate times the number of devices should be a constant. Thus, a cost-balanced chip will have the ratio of chip connections to transistors adjusted equal to the ratio of the transistor to connector, which is equal to 10⁻⁵. This corresponds to the high-end microprocessors that have about 5 million transistors and 500 wire bonds, an identical ratio.

For optoelectronic systems, the defect densities imply that a balanced system will have 10⁴ CMOS transistors per hybrid optoelectronic I/O channel. For a monolithically integrated system, the high electronic defect density implies that the ratio of transistors to optical I/O channels should be ten. This is one reason why monolithic OEICs have been limited to small scales of integration.

To force two balanced systems to have the same cost, the one with a solder-bump connector will have 10 times more I/O channels than one with an hybrid optical connector. To be competitive in performance, the optoelectronic connectors must make up from their lower number with performance advantages.

If the relevant performance metric is bandwidth, a hybrid optoelectronic device may compete by offering 10 times the bandwidth of the solder connector. Since electrical driver power dissipation and wire parasitics limit the electrical bandwidth, avoiding these in optics may allow the necessary 10 times device improvement.

Another way to compensate for the fewer number of hybrid optoelectronic I/Os is to have an architectural advantage. For instance, some graphs like perfect shuffles and hypercubes have large area board layouts that use expensive board area. By using hybrid optoelectronic I/O, the extra board area can be eliminated and the cost reduced.

To balance yield for monolithic integration, the number of transistors must be 1000 times less than the CMOS or hybrid optoelectronic systems. Somehow, these fewer logic devices must give a thousand fold performance advantage to be competitive.

Outside Research Collaborations

The first section of this report on CMOS-SEED manufacturing cost is research in collaboration with Keith Goossen and Jim Walker at Bell Labs. Art D'Asaro, also of Bell Labs, contributed some of the data for the FET-SEED cost model. In addition, Dr. Goossen is working with CU to model the heteroepitaxy process of integration where GaAs is grown directly on silicon. All these model results will be published in a paper that is now in preparation.

To better understand the layout of optical interconnection networks, Mr. Stirk attended a workshop on graph drawing at DIMACS in Princeton. Through a serendipitous encounter, he explained some early work on 3D graph layouts to some attendees working on 3D graphs. This discussion turned into some improved proofs of the volume required for a class of 3D graphs, and is a paper in preparation in collaboration with Peter Eades of University of Newcastle and Sue Whitesides of McGill.

Another recent collaboration is with Marc Desmullicz, John Snowden and Brian Wherret at Heriot-Watt University in Scotland on modeling the cost of European smart pixel manufacturing processes. Together with Heriot-Watt, we put together a survey for the manufacturers to fill out. The survey responses will be used to build cost models for each process. Dr. Desmullicz has applied for a NATO grant to come to CU and continue his research.

Under this grant, Mr. Stirk was invited to attend the Free Space Optoelectronics Workshop on November 4, 1994 at the Air Force Office of Scientific Research. The purpose of the workshop was to chart future application areas of free space optoelectronics. Mr. Stirk participated in the comparisons to electronics, and followed up by working on an optoelectronics technology roadmap.

Technology Transfer

The fundamental investigations of this grant created a basis for several more specific cost modeling efforts. The cost modeling work is an integral part of an ARPA-sponsored effort in collaboration with Hughes Research Labs and UC San Diego that is monitored by RADC/Griffiss. Under this program, we are optimizing the design and process sequence of the Hughes 3D Computer for low-cost manufacturing.

In addition, the cost modeling work is now being applied to a bi-directional optical link program sponsored by ARPA and monitored by Wright Labs. This collaborative effort between AMP, Lasertron, Digital Optics Corporation, Broadband Technologies, CU Boulder and Optoelectronic Data Systems Inc. (ODS) is developing a low-cost transceiver package for fiber in the loop. Mr. Stirk is working at ODS to develop cost modeling software.

Under a separate program sponsored by the Colorado Advanced Technology Institute and ODS, CU is integrating optical, mechanical and thermal modeling CAD/CAM tools with cost tools from ODS. The goal is to create an integrated design and simulation environment for optoelectronics.

Conclusions

The goal of this research grant was a quantitative analysis of the cost and performance of optically interconnected computer architectures. The analysis used manufacturing cost models to compare technologies and system approaches. For the CMOS-SEED process, a hybrid solder bump integration of GaAs modulators onto silicon, the models predict that total yield is dominated by the modulator yield. Though lower-cost than the monolithic FET-SEED, the models predict that the CMOS-SEED chip cost is relatively high due to the expensive GaAs wafer compared to silicon. The high cost and low yield imply that for large systems, the chips will be small to medium scale integration.

By relying solely on device defect rates and balanced system design, we have been able to explain several common microelectronic organizational principles. One principle is the ratio of transistors to I/O pins on a chip. Using the same logic applied to microoptoelectronic technology, we showed that hybrid optoelectronic systems should have 10^4 transistors per optical I/O channel. Monolithic optoelectronic integration should have ten transistors per optical I/O channel.

Bibliography

- C. W. Stirk, Fundamentals of Technology Cost Analysis, invited tutorial, IEEE LEOS Workshop on Interconnects in High-Speed Digital Systems, May 14, 1995.
- C. W. Stirk, Cost-Performance Tradeoffs in Optical Interconnects, OSA Topical Meeting on Optical Computing, Salt Lake, March 13, 1995.
- P. Eades, S. Whitesides, and C. W. Stirk, 3D graph drawing, in preparation.
- C. W. Stirk, K. Goossen, and A. D'Asaro, The cost of optoelectronic integration, in preparation.
- A. R. Mickelson, P. Haugsjaa, C. W. Stirk, M. Stiller, and S. Lee, Results of the 1993 International Workshop on Optoelectronic Packaging Science, submitted to IEEE-CPMT 5/12/94.
- C. W. Stirk, Optoelectronic Integrated Circuit and Multichip Module Manufacturing Cost, submitted to IEEE Trans. on Components, Hybrids, and Manufacturing Technology 2/5/94.
- C. W. Stirk, Introduction to Optical Interconnects, Proc. OCS Student Retreat and Workshop, 47-52, Feb. 18, 1994.
- C. W. Stirk, Parallel Optoelectronic Multiplexer Architectures, First International Workshop on Massively Parallel Processors using Optical Interconnections, April 26-28, 1994.
- C. W. Stirk, Limits on Smart Pixel Granularity: Power Dissipation and Self-Inductance, OSA Proceedings on Photonics in Switching, Vol. 16, 111-114, 1993.

C. W. Stirk, Cost Models of Components for Free-Space Optically-Interconnected Systems, SPIE-1773B-07 (1992).

C. W. Stirk, Hybrid OEIC Cost Sensitivity, LEOS Topical Meeting on Hybrid Optoelectronic Integration and Packaging, 25-26, 1993.